WHAT IS CLAIMED IS:

1. An automatic speech recognition system, comprising:

a speech recognition dictionary comprising a plurality of meaning tokens each associated with one or more pronunciations of one or more vocabulary words and signifying a single meaning; and

a speech recognizer configured to convert spoken input into a sequence of meaning tokens contained in the speech recognition dictionary and corresponding to a sequence of vocabulary words most likely to have been spoken by a user.

2. The system of claim 1, wherein each meaning token is characterized by a unique spelling.

3. The system of claim 2, wherein the spelling of a meaning token facilitates extraction of meaning by a language analyzer.

4. The system of claim 3, wherein the spelling of a meaning token encodes one or more labels identifying one or more respective application-specific categories.

5. The system of claim 4, wherein an application-specific category identified by a label encoded in the spelling of a meaning token is an object category, a place category, an event category, or an action category.

6. The system of claim 1, wherein multiple meaning tokens are associated with each of one or more polysemous vocabulary words contained in the speech recognition dictionary.

7. The system of claim 1, further comprising a language analyzer configured to extract meaning from the sequence of meaning tokens provided by the speech recognizer based upon a set of task-specific semantic rules.

8. The system of claim 7, wherein the language analyzer is a deterministic rule-based language analyzer.

9. The system of claim 7, further comprising an application command translator configured to select an action from a set of application-specific actions based upon the meaning extracted by the language analyzer, and to issue one or more commands to carry out the selected action.

10. The system of claim 1, wherein the speech recognition dictionary is a data structure stored in a computer-readable physical medium.

11. An automatic speech recognition method, comprising:

converting spoken input into a sequence of meaning tokens contained in a speech recognition dictionary and corresponding to a sequence of vocabulary words most likely to have been spoken by a user,

wherein the speech recognition dictionary comprises a plurality of meaning tokens each associated with one or more pronunciations of one or more vocabulary words and signifying a single meaning.

12. The method of claim 11, wherein each meaning token is characterized by a unique spelling.

13. The method of claim 12, wherein the spelling of a meaning token facilitates extraction of meaning by a language analyzer.

14. The method of claim 13, wherein the spelling of a meaning token encodes one or more labels identifying one or more respective application-specific categories.

15. The method of claim 14, wherein an application-specific category identified by a label encoded in the spelling of a meaning token is an object category, a place category, an event category, or an action category.

16. The method of claim 11, wherein multiple meaning tokens are associated with each of one or more polysemous vocabulary words contained in the speech recognition dictionary.

17. The method of claim 11, further comprising extracting meaning from the sequence of meaning tokens based upon a set of task-specific semantic rules.

18. The method of claim 17, further comprising selecting an action from a set of application-specific actions based upon the extracted meaning.

19. The method of claim 18, further comprising issuing one or more commands to carry out the selected action.

20. A computer program for automatically recognizing speech, the computer program residing on a computer-readable medium and comprising computer-readable instructions for causing a computer to:

convert spoken input into a sequence of meaning tokens contained in a speech recognition dictionary and corresponding to a sequence of vocabulary words most likely to have been spoken by a user,

wherein the speech recognition dictionary resides on the computer-readable medium and comprises a plurality of meaning tokens each associated with one or more pronunciations of one or more vocabulary words and signifying a single meaning.